Paper No. 15

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte HANUMA KODAVALLA and NAGAVAMSI PONNEKANTI

Appeal No: 2002-1758
Application No. 09/121,791

DECISION ON APPEAL

Before KRASS, DIXON and SAADAT, Administrative Patent Judges.

KRASS, Administrative Patent Judge.

DECISION ON APPEAL

This is a decision on appeal from the final rejectionof claims 1-30, all of the pending claims.

The invention is directed to a database system providing for high-concurrency access in B-Tree structures.

Representative independent claim 1 is reproduced as follows:

> 1. In a database system, said database system storing a plurality of data records as a data table having an index, said index

comprising a B-Tree having a root page, a plurality of internal pages, and a plurality of leaf-level pages, each page storing one or more key values from said data records, a method for providing highly-concurrent access to the B-Tree for multiple clients, the method of comprising:

(a) receiving a request from a client which requires insertion of a new key value into the B-Tree in order to satisfy the request;

(b) traversing the B-Tree for locating an existing page in the B-Tree appropriate for storing the new key value; and

(c) if said existing page does not contain sufficient room for storing the new key value, splitting the existing page into two pages by:

(i) allocating a new page at a level in the B-Tree which is the same as the existing page and marking both pages as undergoing a split,

(ii) moving some of the key values from the existing page to the new page, and

(iii) creating a reference in the existing page which points to the new page, so that any client which is traversing the B-Tree will not be blocked by the split which is currently occurring.

The examiner relies on the following references:

Ishak et al. (Ishak)	5,475,837	Dec. 12, 1995
Roy	5,644,763	Jul. 1, 1997

Claims 1-30 stand rejected under 35 U.S.C. §103 as unpatentable over Roy in view of Ishak.

Reference is made to the brief and answer for the respective positions of appellants and the examiner.

OPINION

We begin our analysis with a definition of B-tree structures. Since neither appellants nor the examiner define this term, and its meaning is not precisely known from the applied references, we look at the definition supplied by the National Institute of Standards and Technology (NIST):

> A *balanced search tree* in which every *node* has between [m/2] and m *children*, where m>1 is a fixed integer. M is the *order*. The *root* may as few as 2 children [sic]. This is a good structure if much of the tree is in slow memory (disk), since the *height*, and hence the number of accesses, can be kept small, say one or two, by picking a large m. [Google internet search on March 2, 2004, at http://www.nist.gov/dads/HTML/btree.html].

Appellants seek to provide high-concurrency access to B-tree structures. By way of explanation, page 4 of the instant specification indicates that when the system receives a request to insert a key value into a B-tree at a page that does not have sufficient room, the system must split at the tree at leaf level. This is performed by allocating a new page and moving some of the key values from the old page to the new page. The specification explains that the split itself propagates upward and to do the split itself, the system must obtain address locks for the two pages, marking both as undergoing "split". The system would then add the address locks as a linked list of address locks. When the split is propagated up, a "side entry" is added to the old page to point to the newly allocated page and, since the old page may not have sufficient room

for storing the new entry, the parent page must split also. The parent page split is performed by allocating a new page at that level and then both pages, old and new parents, are marked as undergoing split. The system still obtains address locks for these pages but, at this point, a side entry is created in the old parent page. This information lets any client searching for a key value know that, instead of going directly down the tree from the old parent page, it should proceed to the parent's new sibling node or page.

As explained, at page 5 of the specification,

> ...the traversal is not blocked by the split which is currently active. In effect, the client knows how to take a detour to arrive at the proper leaf-level page. Similarly, if the client is searching for a lesser key value, it will traverse to the proper page, without being blocked. Again, even though the page is undergoing a split, the client is not blocked from accessing the sought-after key value. After split propagation is complete, the system clears the split flags and releases address locks. Also, at this point, the side entry is removed. Now, sufficient room now exists in the tree for inserting the key value.

The key elements of allocating a new page "at a level in the B-Tree which is the same as the existing page and marking both pages as undergoing a split" and "creating a reference in the existing page which points to the new page, so that any client which is traversing the B-Tree will not be blocked by the split which is currently occurring" are clearly set forth in instant independent claim 1. Independent claim 21 clearly sets forth that the means for inserting a new key value includes "means for splitting the existing page into the existing page and a new page, said new page being at a level in the B-Tree which is the same as the existing page" and claim 21

further sets forth that the means for preserving concurrent access includes "a side-entry storing a key value allowing a client traversing the B-Tree to determine whether to traverse to the new page."

With regard to independent claim 1, the examiner asserts that

> Roy and Ishak teach the invention substantially as claimed. Ishak further teaches (I) allocating a new page at a level in the B-Tree which is the same as the existing page and marking both pages as undergoing a split and moving some of the key values from the existing page to the new page [702 of fig. 7]. [Paper No. 5-page 3].

With regard to independent claim 2, the examiner asserts that Roy discloses the claimed invention but for an explicit teaching of

> means for preserving concurrent access by creating an entry in the existing page which points to the new page, so that any other client which is traversing the B-Tree at the point of the split will not be blocked by the split while it is occurring although it has the same functionality of splitting pages at insertion points [see the abstract]. However, Ishak teaches structural modification processes such as a page splitting can be carried out on a page which is concurrently being accessed by any other client [ab; 704, 706 of fig. 7; col. 4, lines 29-44]. [Paper No. 8-pages 3-4].

The examiner then concludes that it would have been obvious to add Ishak's features to Roy "in order to increase access and performance in B-tree structures in which the traversal is not blocked by the split which is currently active." [Paper No. 8-page 4].

We have carefully reviewed the evidence in the case, including the arguments of both appellants and the examiner and we conclude that appellants make compelling arguments which, in our view, distinguish the instant claimed invention from that taught by the applied references, and which have not been satisfactorily answered by the examiner.

For example, both of the instant independent claims require the allocation of the new page *at the same level* as the existing page. While there is no doubt that Ishak (relied on by the examiner for page splitting on a page concurrently being accessed) provides an approach for maintaining concurrency, it does not appear to describe the instant claimed node splitting mechanism (i.e., split bit and side entry) in which a newly-allocated page or node is temporarily linked at the same level as the pre-existing page that is being split.

The instant claims also provide for a side-entry link and accompanying split bits, identified by appellants as "marking both pages as undergoing a split," for referencing a sibling node being created, in order to support highly-concurrent traversal of the B-Tree while the split is occurring (see page 8 of the brief). Appellants distinguish this from Ishak's conventional approach of updating a parent node to point to newly-created children nodes, with Ishak specifically creating two new nodes.

Appellants further contrast Ishak's step 706 in Figure 7, requiring updating the old link in the parent node to point to the HI node instead of the old node, with the instant claims, requiring split bits and side-entry links when creating a new node at the same level (i.e., a sibling node) as

the existing node. Whereas Ishak describes updating a parent node to point to children nodes, the instant claims describe a node or page temporarily creating a linkage to a sibling node being created [see page 9 of the brief].

The examiner has not satisfactorily responded to appellants' rather convincing arguments. At pages 3-4 of the answer, the examiner contends that appellants' argued distinctions are based upon "features...not recited in the rejected claims." We disagree. As appellants explain, the claim limitation of "marking both pages as undergoing a split," as set forth in claim 1, refers to the disclosed side-entry link and accompanying split bits (brief-page 8). Reference to the specification supports appellants' explanation and such reference is permitted in order to ascribe a meaning to a claimed term. Both independent claims 1 and 21 recite the allocation of a new page at the same level (i.e., sibling node) as the existing page. Claim 21 is explicit in its recitation of a "side-entry" storing a key value.

The examiner also takes issue with appellants' argument regarding allocation of a new page at the same level as the existing page. In particular, the examiner points to Roy's description of the splitting of a single node by allocating a new node, in Figure 3, and to Ishak's linking to new nodes in the splitting process. "It is clear from the fig. 2-6B that links at the same level exist" (answer-page 4).

We are not quite sure of the point the examiner is trying to make. However, it is clear to us that Ishak is concerned with updating a parent node to point to children nodes, and not to

sibling nodes, which is what the "same level" language in the instant claims concerns. Since Roy is not concerned with preserving concurrent access by creating an entry in the existing page which

points to the new page in the manner described in the instant claims, it is not clear to us how any modification by Ishak would result in the claimed subject matter wherein the allocation of a new page is at the same level as the existing page.

The examiner has not convinced us that the combination of Roy and Ishak presents substantial evidence adequate for holding the instant claimed subject matter obvious in accordance with 35 U.S.C. §103.

Accordingly, the examiner's decision rejecting claims 1-30 under 35 U.S.C. §103 is reversed.

REVERSED

ERROL A. KRASS Administrative Patent Judge	)	
JOSEPH L. DIXON Administrative Patent Judge	)	BOARD OF PATENT APPEALS AND INTERFERENCES

)
)
MAHSHID D. SAADAT)
Administrative Patent Judge)

EAK/dpv

JOHN A. SMART
SYBASE, INC.
5000 HACIENDA DR.
DUBLIN, CA 94568